arXiv:1605.03391v1 [stat.ML] 11 May 2016

# Random forests for survival analysis using maximally selected rank statistics

Marvin N. Wright[1], Theresa Dankowski[1] and Andreas Ziegler[1–4*]

Feb. 19, 2016

**Abstract**

The most popular approach for analyzing survival data is the Cox regression model. The Cox model may, however, be misspecified, and its proportionality assumption is not always fulfilled. An alternative approach is random forests for survival outcomes. The standard split criterion for random survival forests is the log-rank test statistics, which favors splitting variables with many possible split points. Conditional inference forests avoid this split point selection bias. However, linear rank statistics are utilized in current software for conditional inference forests to select the optimal splitting variable, which cannot detect non-linear effects in the independent variables. We therefore use maximally selected rank statistics for split point selection in random forests for survival analysis. As in conditional inference forests, p-values for association between split points and survival time are minimized. We describe several p-value approximations and the implementation of the proposed random forest approach. A simulation study demonstrates that unbiased split point selection is possible. However, there is a trade-off between unbiased split point selection and runtime. In benchmark studies of prediction performance on simulated and real datasets the new method performs better than random survival forests if informative dichotomous variables are combined with uninformative variables with more categories and better than conditional inference forests if non-linear covariate effects are included. In a runtime comparison the method proves to be computationally faster than both alternatives, if a simple p-value approximation is used.

## 1 Introduction

Survival data are generally analyzed using regression models, such as the Cox proportional hazards model [1]. The standard Cox model is not well suited for big data problems, such as large scale omics studies or image analysis. Regularized Cox models may be used in this case [2]. However, the model may be misspecified, and the assumption of proportional hazards may also be violated. A promising alternative is the use of machine learning methods, such as survival trees or random forests [3], and several tree-based approaches have been proposed for survival outcome [4–7]. The two major algorithms for random forests with survival outcome are random survival forests [8, RSF] and conditional inference forests [9, CIF]. Recently, the statistical properties of random forests, including random survival forests have been better understood. Specifically, results have been obtained on consistency, convergence rate and asymptotic normality of random forest estimators [10–13].

RSF are closely related to the original random forests approach by Breiman [3]. Specifically, at each node the best split is selected as the maximum of the log rank statistic over all possible split points over all independent variables available for splitting at that node. In the standard R package for RSF

*E-mail: ziegler@imbs.uni-luebeck.de

[1]Institut für Medizinische Biometrie und Statistik, Universität zu Lübeck, Universitätsklinikum Schleswig-Holstein, Campus Lübeck, Lübeck, Germany

[2]Zentrum für Klinische Studien, Universität zu Lübeck, Lübeck, Germany

[3]School of Mathematics, Statistics and Computer Science, University of KwaZulu-Natal, Pietermaritzburg, South Africa

[4]Deutsches Zentrum für Herz-Kreislauf-Forschung, Standort Hamburg/Kiel/Lübeck, Lübeck, Germany

randomForestSRC [14], tree growing is stopped if the terminal node size is below a pre-defined threshold. A disadvantage of this random forests approach is that it favors splits for covariates with many possible split points [15–18]. For illustration, consider a dataset with response variable $y$ and two covariates $x_1$ and $x_2$ with $n_1$ and $n_2$ possible cut points, respectively. Further, suppose that $y$ is independent of both $x_1$ and $x_2$, and $n_1 < n_2$. If we now search for the best split point by comparing all possible split points of both independent variables for their effect on $y$, $x_2$ will have a higher probability to have the split point with the larger effect on $y$—just by chance. If the split point selection is biased, other parameter estimates, such as variable importance measures are biased as well [17]. Predictions could also suffer from the underestimation of important variables with few categories. In particular, if gene expression data and single nucleotide polymorphisms (SNP), both measured on microarrays, are to be analyzed jointly, the quantitative gene expression levels will be generally preferred for splitting by the RSF algorithm over the discrete SNP data. In fact, large-scale microarray data contain many uninformative, i.e., noise variables, and they are potentially mixed with important clinical covariables. This increases the need for unbiased split variable selection [19].

An approach to reduce split point selection bias are CIF. CIF are conceptually different from standard random forests because they separate the selection of the split variable from selection of the split point of the already selected split variable [9]. In the first step, the optimal split variable is determined. Specifically, an association test is performed between the possible split variables and the response. In the party package [20], this association test is a linear rank test. For survival data, a log-rank transformation for censored data is performed. If the association is found to be significant, the covariate with minimal $p$-value is selected for splitting. If no significant association is found, no split is conducted. In the second step, the optimal split point is found by comparing two-sample linear statistics for all possible partitions for the split variable. The two-step approach of CIF leads to unbiased split point selection if sampling is done without replacement and if a test statistic of quadratic form is used [17]. However, the selection bias is ignored by many researchers using random forests [21].

The disadvantage of currently implemented CIF is that the association test for selecting the split variable is based on a linear rank statistic, while the optimal split is a dichotomous threshold-based split. In this work, we propose to use maximally selected rank statistics [22–24] for CIF to avoid a change in the statistical approach for split variable and split point selection. With maximally selected rank statistics, the optimal split variable is determined using a statistical test for binary splits, which adjusts for the multiple testing of multiple possible split points. Thus, this approach naturally avoids split point selection bias. For example, for dichotomous split variables, such as sex, no adjustments for multiple testing are required. In contrast, through the use of maximally selected rank statistics adjusted $p$-values are obtained for continuous covariates, such as gene expression levels. Since the exact distribution of the maximally selected rank statistic is generally unknown, $p$-values need to be approximated. We have implemented maximally selected rank statistics for random forests for survival endpoints using several $p$-value approximations. We have investigated the split point selection, prediction performance and runtimes of the different approaches in simulation studies and with real data.

In the next section, maximally selected rank statistics for survival outcome are introduced. Their implementation in a survival RF algorithm is described in Section 3. In Section 4, the selection bias of different methods for the $p$-value approximation is investigated, and in Section 5 the prediction performance of the new approach is compared with existing methods on simulated data. Finally, in Section 6, prediction performance and runtime of the methods is evaluated on three real datasets on breast cancer, including a gene expression study and a genome-wide association study.

## 2 Maximally selected rank statistics for survival endpoints

In this section, we aim at finding an optimal cutpoint to split data in two groups using maximally selected rank statistics. A cutpoint is considered optimal if the separation of the survival curves in the two groups is maximized.

We consider $n$ observations $(X_1, Y_1), \ldots, (X_n, Y_n)$ and a continuous marginal distribution for $X$. $Y_i = (Z_i, \delta_i)$ is a right censored survival time with time $Z_i$ and censoring indicator $\delta_i$. We use log-

rank scores $a_1, \dots, a_n$ [25]

$$a_i = a_i(\mathbf{Z}, \boldsymbol{\delta}) = \delta_i - \sum_{j=1}^{\gamma_i(\mathbf{Z})} \frac{\delta_j}{(n - \gamma_j(\mathbf{Z}) + 1)},$$

where $\mathbf{Z} = (Z_1, \dots, Z_n)'$ and $\boldsymbol{\delta} = (\delta_1, \dots, \delta_n)'$ are the vectors of the survival time and censoring indicator, respectively, and $\gamma_j(\mathbf{Z}) = \sum_{i=1}^{n} \mathbb{1}_{\{Z_i \leq Z_j\}}$ is the number of observations with survival time up to $Z_j$. The linear rank statistic for a split by a cutpoint $\mu$ is the sum of all scores in the group with $X_i \leq \mu$

$$S_{n\mu} = \sum_{i=1}^{n} \mathbb{1}_{\{X_i \leq \mu\}} a_i.$$

The null hypothesis of no influence of a split by the cutpoint $\mu$ on the distribution of $Y$ is $H_0 \colon P(Y \leq y | X \leq \mu) = P(Y \leq y | X > \mu)$ for all $\mu$ and all $y$. Under the null hypothesis the expectation and variance are given by [22]

$$\mathbb{E}_{H_0}(S_{n\mu} | a, X) = m_\mu \bar{a} \quad \text{and} \quad \mathbb{V}ar_{H_0}(S_{n\mu} | a, X) = \frac{n}{n-1} \frac{m_\mu}{n} \frac{n_\mu}{n} \sum_{i=1}^{n} (a_i - \bar{a})^2,$$

where $m_\mu = \sum_{i=1}^{n} \mathbb{1}_{\{X_i \leq \mu\}}$ and $n_\mu = \sum_{i=1}^{n} \mathbb{1}_{\{X_i > \mu\}} = n - m_\mu$ are the number of observations in the two groups, and $\bar{a} = \frac{1}{n} \sum_{i=1}^{n} a_i$ is the mean of the scores of all observations. To compare different splits we use the score test statistic

$$T_{n\mu} = \frac{S_{n\mu} - \mathbb{E}_{H_0}(S_{n\mu} | a, X)}{\sqrt{\mathbb{V}ar_{H_0}(S_{n\mu} | a, X)}}.$$

To maximize the separation of the data induced by a split, the cutpoint $\mu$ with maximal $T_{n\mu}$ is selected. To guarantee sufficiently large sample size in each group, the possible cutpoints $\mu$ are restricted to $\mu_1 \leq \mu \leq \mu_2$ with $\mu_1$ and $\mu_2$ corresponding to quantiles $\varepsilon_1, \varepsilon_2$ of the distribution of $X$. The maximally selected rank statistic is defined as

$$M_n(a, X, \varepsilon_1, \varepsilon_2) = \max_{\mu \in [\mu_1, \mu_2]} (|T_{n\mu}|).$$

In standard random forests the values of $T_{n\mu}$ would be compared not only between split points on the same variable but also between variables. As explained in Section 1, this would induce a bias if the number of possible splitting points differs between variables. Therefore, we obtain $p$-values for the maximally selected rank statistic $M_n(a, X, \varepsilon_1, \varepsilon_2)$ for each variable and compare the variables on the $p$-value scale. However, the exact distribution $P_{H_0}(M_n(a, X, \varepsilon_1, \varepsilon_2) \leq b)$, $b \in \mathbb{R}^+$ of the test statistic is generally unknown, and no exact $p$-values can be obtained. Several approximations to the distribution of the exact test statistic under the null hypothesis have been proposed, and asymptotic $p$-value approximations can be obtained as well. First, the exact distribution can be approximated using a standard permutation approach. Second, to approximate the $p$-value, Lausen and Schumacher [22] proposed using the distribution of the supremum of the absolute value of a standardized Brownian bridge. Lausen *et al.* [23] provided an approximation using an improved Bonferroni inequality [26]. Lausen *et al.* [23] recommended to use the minimum $p$-value of the two approximation methods because both are conservative. Hothorn and Lausen [25] derived a lower bound of the distribution of the test statistic using the exact distribution of linear rank statistics by utilizing Streitberg and Röhmel's shift algorithm [27]. Finally, the distribution of the test statistic can be approximated by a maximally selected Gauss statistic [25] because linear rank statistics are asymptotically normally distributed under suitable regularity conditions [28]. The distribution of a maximally selected Gauss statistic can be computed using the algorithms by Genz [29].

# 3 Survival forests with maximally selected rank statistics

In this section, we propose a random forest algorithm for survival endpoints using maximally selected rank statistics (MSR-RF), compare it with standard random forests, RSF and CIF, describe the estimation of specific parameters and the implementation of the procedure using R and C++.

### 3.1 Core elements of the algorithm

The main elements of the algorithm are identical to those of the standard random forest algorithm. The number of bootstrap samples is fixed a priori. Next, a bootstrap sample is drawn from the data, and a tree is grown for each bootstrap sample. Alternatively, subsampling of the observations can be employed instead of bootstrapping. To induce randomness on the covariate level, a subset of covariates is randomly selected and made available for splitting at each node. The tree growing is stopped if a specific stop criterion is fulfilled. In the final step of the algorithm, results of the trees are aggregated, e.g., to estimate a prediction error.

Our novel algorithm follows the basic concept of CIF. Thus, we propose to split nodes using a two-step procedure. First, for each candidate covariate, maximally selected rank statistics are computed, and the split point with the maximal standardized test statistic is selected. For each of these covariates, the $p$-value is obtained for the best split point under the null hypothesis of no association between the split point and the covariate. $p$-values are approximated or asymptotic $p$-values, estimated using one of the methods described in Section 2. The covariate with the smallest $p$-value is selected as splitting candidate, where $p$-values are adjusted using the Benjamini and Hochberg [30] adjustment for the multiple testing of a total of `mtry` variables at a possible split point. If the adjusted $p$-value of the candidate covariate is not smaller than a pre-specified type I error level $\alpha$, no split is performed. Thus, only if the adjusted $p$-value of the candidate covariate is smaller than $\alpha$, the split is made. This part is one of the major differences between the proposed random forest algorithm and the standard random forest procedure. A major difference to CIF is in the second step: In CIF, the optimal split point is determined here. This simplifies in the proposed approach because the optimal split point is determined as a by-product in step 1 of the procedure. It is the split point corresponding to the maximally selected rank statistic in step 1.

In the original CIF approach [9], a linear rank test between log rank scores and the covariate is implemented in `party`. Thus, a linear association test is performed between a possibly continuous covariate and quantitative score values. Although a standard association test is performed in the first step, a standard binary split is done in the second step. With the MSR-RF procedure, the optimal binary split is determined as in standard random forests, but an adjustment for multiple possible splits is performed through the use of maximally selected rank statistics. As a result, no adjustments need to be made for dichotomous variables, such as sex. However, several adjustments are made for ordinal covariates, and adjustments play a major role for continuous covariates. An advantage of MSR-RF when compared with CIF as implemented in `party` is that binary splits are used for determining the best split and the optimal splitting variable. In this way, one procedure is used consistently in both steps of MSR-RF.

### 3.2 Estimation

In each terminal node $v$ of a tree, the survival function is estimated using the Kaplan-Meier estimator, utilizing only the observations from the same terminal node. For a single tree $b$, a prediction of the survival function $\hat{S}_b(t|\boldsymbol{x}_i)$ is made for a new observation $\boldsymbol{x}_i$ by dropping the observation down the tree. The prediction of the random forest $f$ with $B$ trees is obtained by averaging the predictions over all trees for each observation and time point:

$$\hat{S}_f(t|\boldsymbol{x}_i) = \frac{1}{B}\sum_{b=1}^{B} \hat{S}_b(t|\boldsymbol{x}_i). \tag{1}$$

It is standard in random forests to use only the out-of-bag (OOB) samples for estimating the prediction error. We define $I_{i,b} = 1$ if $i$ is an OOB observation in tree $b$, and 0, otherwise [8]. The OOB prediction for observation $i$ is then defined as

$$\hat{S}_f^{OOB}(t|\boldsymbol{x}_i) = \sum\nolimits_{b=1}^{B} \left(I_{i,b} \cdot \hat{S}_b(t|\boldsymbol{x}_i)\right) \Big/ \sum\nolimits_{b=1}^{B} I_{i,b}\,. \tag{2}$$

The prediction error is estimated with the Brier score (BS) for censored data [31] using only the

out-of-bag observations of each tree

$$BS(t, \hat{S}_f^{OOB}) = \frac{1}{n} \sum_{i=1}^{n} \widehat{W}_i(t) \left( \mathbb{1}_{\{T_i > t\}} - \hat{S}_f^{OOB}(t|\boldsymbol{x}_i) \right)^2,$$

where $T_i$ is the observed survival time for subject $i$. $\hat{S}_f^{OOB}(t|x_i)$ is the random forest estimator of the survival function, and $\widehat{W}_i(t)$ are weights using the inverse probability of censoring [32]. For $t \in [0, t^*]$, the integrated Brier score ($IBS$) is obtained as

$$IBS(\hat{S}_f^{OOB}) = \int_0^{t^*} BS(t, \hat{S}_f^{OOB}) \, \mathrm{d}t.$$

### 3.3 Implementation in R

The new method is implemented in the R [33] package `maxstatRF` which is supplied with this article. To avoid unnecessary copying of data, the package is implemented using reference classes [34]. Furthermore, the `parallel` package [33] is used for parallel computation, and the package `ipred` [35] is employed for calculating the BS. The `maxstat` package [36] provides all necessary elements for the computation of maximally selected rank statistics. In detail, the following $p$-value approximations are used: Lausen and Schumacher [22, `Lau92`], Lausen *et al.* [23, `Lau94`], Hothorn and Lausen [25, `HL`], permutation (`condMC`) and the approximation by the maximally selected Gauss statistic (`exactGauss`). Finally, the minimum of `Lau92` and `Lau94` was added (`minLau`).

Tuning parameters in the `maxstatRF` package include standard random forest settings, such as the number of trees (`num_trees`) in the random forest, the number of randomly selected variables considered for splitting in each node (`mtry`) and the minimal terminal node size (`min_node_size`). The parameter `minprop` corresponds to the lower quantile $\varepsilon_1$ of the covariate distribution (see Section 2). The upper quantile is always set to $\varepsilon_2 = 1 - \varepsilon_1$. The significance level for splitting $\alpha$ and the $p$-value approximation method can be set as in the `maxstat` package. Finally, the sampling can be done with replacement (bootstrapping) or without replacement (subsampling). If the latter is used, the bootstrap sampling in the RF algorithm is replaced by subsampling with approximately 63.2% of the observations.

### 3.4 Implementation in C++

The software package `maxstatRF` is not optimized for computational efficiency. We also implemented the new method in the R package `ranger` [37] to enable the efficient analysis of larger datasets, the growth of random forests with large `mtry` values, the tuning of the terminal node size or a large number of trees. The results from Section 5 show that the $p$-value approximation by the minimum of the approximations by Lausen and Schumacher [22] and Lausen *et al.* [23] leads to high prediction accuracy. Since other approximation methods show lower prediction accuracy or are computationally more intensive (Section 4), only this approximation is implemented in `ranger`.

## 4 Simulation study 1: split point selection

Since approximations are used for the distribution of the maximally selected rank statistic, the split point selection is potentially biased. Furthermore, the use of bootstrapping could induce a biased selection [17]. To assess both sources of potential bias we conducted a simulation study. Datasets were generated with a censored survival outcome and 5 covariates. The survival outcome was simulated by the minimum of a survival and censoring time, both assumed to be exponentially distributed with $\lambda = 0.5$ and $\lambda = 0.1$, respectively. Four out of 5 covariates were simulated with a multinomial distribution with a varying number of categories: $x_2$, $x_4$, $x_{10}$ and $x_{20}$ with 2, 4, 10 and 20 categories, respectively, and equal probabilities for each category. The $5^{th}$ covariate $x_n$ was simulated using a standard normal distribution to generate many unique values for the dataset. We simulated the null model [17] of totally uninformative predictor variables. In the ideal case, all covariates should thus be selected with equal frequencies.

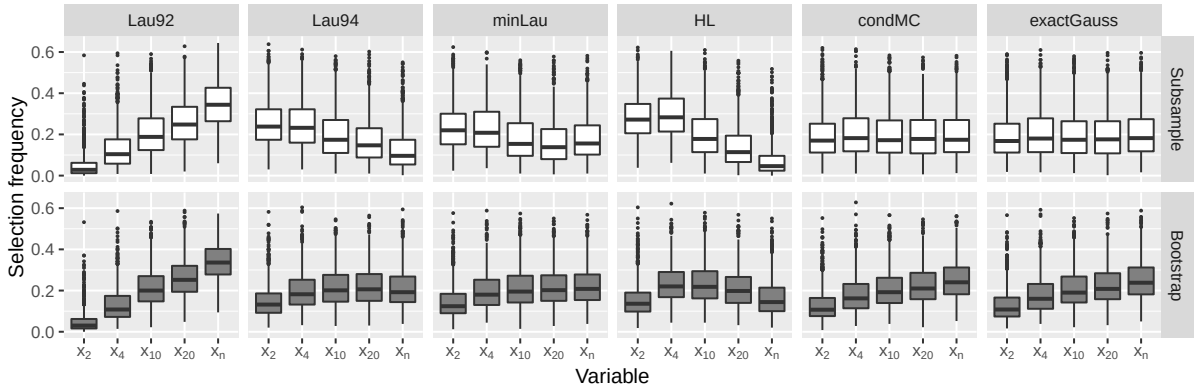


Figure 1: Split point selection for different $p$-value approximation methods for the null case of no association between covariates and survival time. The $p$-value approximations of Lausen and Schumacher [22, `Lau92`], Lausen *et al.* [23, `Lau94`], their minimum (`minLau`), the approximation by Hothorn and Lausen [25, `HL`], the permutation approach (`condMC`) and the maximally selected Gauss statistic (`exactGauss`) are compared. Subsampling, i.e., sampling without replacement, is displayed in the upper panel and bootstrapping, i.e., sampling with replacement, in the lower panel. The covariates $x_2$, $x_4$, $x_{10}$ and $x_{20}$ were simulated with 2, 4, 10 and 20 categories, respectively, and the covariate $x_n$ is continuous.

All simulations were run with 500 trees, `mtry` $= 3$ and covariate quantiles of $\varepsilon_1 = 0.1$, thus $\varepsilon_2 = 0.9$. All 6 $p$-value approximation methods were compared. Only the first split was investigated in each tree to avoid the influence of previous splits. For example, after one split on a dichotomous variable in a tree, this variable could not be selected again in that branch, and no unbiased split point selection could be measured if further splits would be considered. We furthermore analyzed the effect of sampling with replacement (bootstrapping) and without replacement (subsampling). For subsampling a proportion of 63.2% of the observations was used, which approximately equals the number of unique observation in a bootstrap sample. The number of replications was set to 1000.

Fig. 1 displays the results for the different approximations at a sample size of $n = 100$. The Lausen and Schumacher `Lau92` approximation showed preferential splitting of covariates with a large number of possible split points for both bootstrapping and subsampling. With Lausen *et al.* (`Lau94`), variables with few categories were preferred if subsampling was used and variables with many categories for bootstrapping. However, the bias was less pronounced than for the Lausen and Schumacher approximation. The combination `minLau` was similar to the approximation by Lausen et al., but in the case of subsampling the bias was slightly lower. The Hothorn and Lausen approximation `HL` showed selection bias towards covariates with few categories for both resampling schemes, and it was quite large in case of subsampling. Both the permutation approach `condMC` and the exact Gauss statistic yielded unbiased splits in case of subsampling and preferential splitting of covariates with many categories in case of bootstrapping.

Supplementary Figures $S1 - S2$ display the results for varying sample sizes of $n = 25, 50, 100$ and $n = 200$ and covariate quantiles varying between $\varepsilon_1 = 0.1, 0.25$ and $0.4$. No results were obtained for the Hothorn and Lausen $p$-value approximation for $n = 200$ and bootstrapping because the approximation cannot handle sample sizes $> 171$. More specifically, the gamma function $\Gamma(\cdot)$ is used to compute the number of permutation, and $\Gamma(n)$ is out of double precision range on a 64 bit computer system if $n > 171$. Otherwise, results were similar to those displayed in Figure 1.

The split point selection of RSF and CIF was investigated (Figure $S3$, supplementary material). As expected from the results of Strobl *et al.* [17], RSF was biased towards variables with many categories for subsampling and bootstrapping, while CIF was unbiased for subsampling and biased for bootstrapping.

Since the $p$-value approximation is computed for each of the `mtry` splitting candidate covariates at every possible split in each tree, growing a forest with maximally selected rank statistics can be computationally intensive. We therefore compared the runtimes of the approximation methods for varying sample sizes. Data was simulated as in the previous simulation study with samples sizes between 10 and

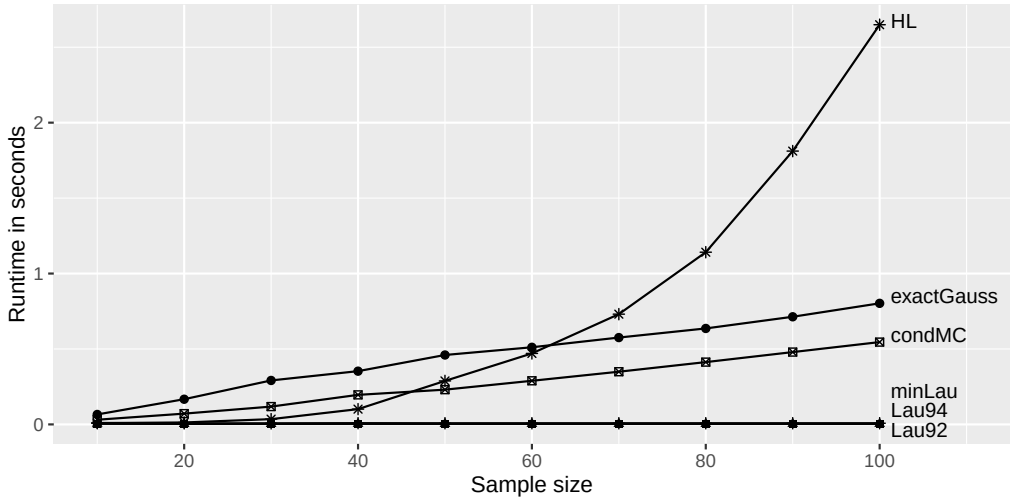


Figure 2: Runtime (in seconds) for varying sample size for different $p$-value approximations. Displayed runtimes correspond to the computation required in each node of a tree if 5 covariates are considered and data is simulated as in Figure 1. The $p$-value approximations of Lausen and Schumacher [22, `Lau92`], Lausen *et al.* [23, `Lau94`], their minimum (`minLau`), the approximation by Hothorn and Lausen [25, `HL`], the permutation Monte Carlo approach (`condMC`) and the maximally selected Gauss statistic (`exactGauss`) are compared. The runtimes of `Lau92`, `Lau94` are almost identical. The runtime of `minLau` is the sum of `Lau92` and `Lau94`.

100, and the $p$-values of the covariates were approximated using the different methods. The runtime results are summarized in Figure 2 and Table S1 (supplementary material). The runtimes of `Lau92` and `Lau94` were approximately equal and smaller than for all other approximations. The runtime of the `minLau` method is the sum of the runtimes of `Lau92` and `Lau94`. The `HL` approximation was fast for small sample sizes, but runtime grew exponentially for larger sample sizes. Both `condMC` and `exactGauss` scaled approximately linearly with the number of sample sizes, with a larger slope for `exactGauss`.

# 5 Simulation study 2: prediction performance

To assess the prediction performance of the new method with the different $p$-value approximations, we conducted a second simulation study. We simulated datasets of two simulation scenarios and analyzed them with MSR-RF and all 6 $p$-value approximations, RSF, CIF and the Cox model.

The first simulation scenario was specifically designed to have few highly predictive variables with few categories and many non-predictive variables with many categories. The scenario was simulated with 200 subjects, 6 dichotomous covariates with strong effect and 100 normally distributed independent variables without effect. The second scenario consisted of 200 subjects and 14 covariates. Of these, 2 variables were continuously uniformly distributed with a strong non-linear effect as depicted in Figure 3. Two covariates were Bernoulli distributed and had a moderate linear effect, and 10 covariates were normally distributed without effect. In both scenarios, the survival times and censoring times were simulated with exponential distributions.

To avoid overfitting, 10-fold cross-validation was employed to assess prediction accuracy. For each cross-validation fold, a model was built on the remaining 9 subsamples, and the prediction error was estimated on the selected fold. The integrated Brier score for censored data [31, IBS] was used for estimating the prediction errors, see Section 3. The R package `pec` [32] was used for computing the IBS. All random forest analyses were run using 50 trees and `mtry` $= p/2$, where $p$ is the number of covariates in the model. For MSR-RF, $\varepsilon_1 = 0.1$, and all 6 $p$-value approximation methods were used. All

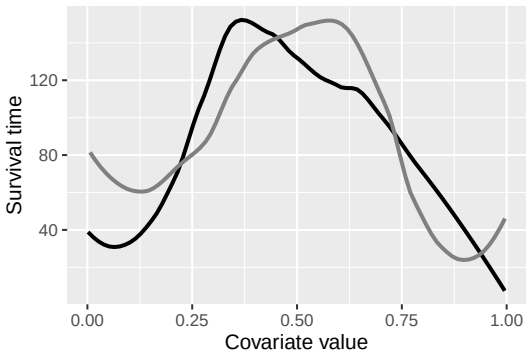


Figure 3: Effects of the non-linear covariates in simulation scenario 2. The covariates were drawn from a continuous uniform distribution $U(0, 1)$. The plotted data is LOESS smoothed with a span of 0.5.

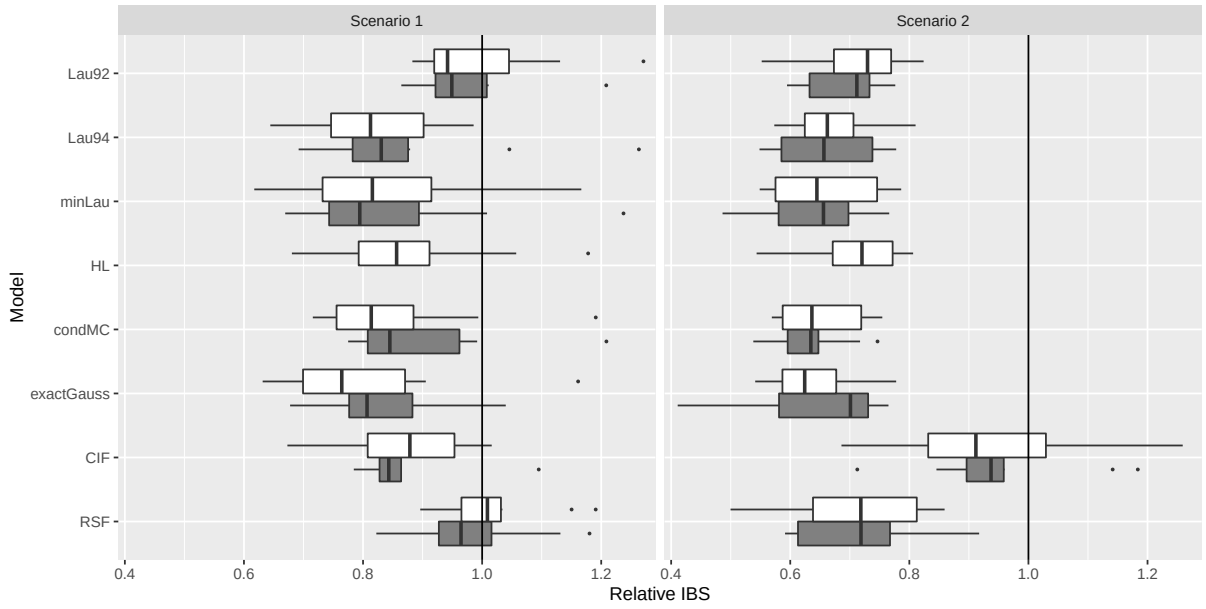


Figure 4: Prediction error for two simulation scenarios, different analysis methods and resampling schemes. The prediction errors are shown by boxplots (median, quartiles, largest non outliers) of integrated Brier scores (IBS) relative to the Cox model. Values left to the vertical line show results with lower IBS than the Cox model. White boxed correspond to subsampling, where 63.2% of the samples were resampled without replacement. Grey boxes correspond to bootstrapping. The $p$-value approximations in MSR-RF are Lausen and Schumacher [22, `Lau92`], Lausen *et al.* [23, `Lau94`], their minimum (`minLau`), the approximation by Hothorn and Lausen [25, `HL`], the permutation Monte Carlo approach (`condMC`) and the maximally selected Gauss statistic (`exactGauss`). In addition, random survival forests [8, RSF] and conditional inference forests [9, CIF] are displayed. No results for `HL` and bootstrapping could be obtained because the approximation cannot handle sample sizes $> 171$; see Section 4 for details.

analyses were repeated for sampling with replacement (bootstrapping) and sampling without replacement (subsampling). In each run, all models were trained on the same data.

As described in Section 3, the size of the trees in MSR-RF is determined by the significance level for splitting $\alpha$. For small values the tree growing is stopped early, while larger values of $\alpha$ result in larger trees. In CIF, the tree size is also determined by a significance level, but a parameter `mincriterion` $= 1 - \alpha$ is used. In RSF, no significance test is performed for splitting, and the tree size is controlled by limiting the minimal terminal node size with a parameter `nodesize`. The optimal size of trees heavily depends on the algorithm, the resampling method and the data. Therefore, `nodesize` was tuned in a preceding round of cross-validation for each algorithm, dataset and subsampling combination. The parameters $\alpha = 0.1, 0.3, 0.5, 0.7, 0.9$, `nodesize` $= 1, 3, 10, 25, 50$ and `mincriterion` $= 0.1, 0.3, 0.5, 0.7, 0.9$ were evaluated for MSR-RF, RSF and CIF, respectively. For each combination, the parameter was selected that led to the lowest average IBS on all cross-validation folds. Different random seeds were set for tuning and final comparison. The optimal parameters are displayed in Table S2 (supplementary material).

The results of the benchmark experiments are shown in Figure 4. In the first simulation scenario (Scenario 1), MSR-RF performed best with the `Lau94`, `minLau` and `exactGauss` approximations, followed by the `condMC`, `HL` and CIF, with only small differences. MSR-RF with the `Lau92` approximation and RSF performed similarly to the Cox model, but worse than the other methods. No systematic difference between subsampling and bootstrapping could be observed. In the second simulation scenario (Scenario 2), MSR-RF performed best with the `Lau94`, `minLau`, `condMC` and `exactGauss` approximations, followed by RSF, `Lau92` and `HL`. CIF and the Cox model performed substantially worse than the other models. Again, no systematic difference could be observed between subsampling and bootstrapping.

As expected, the RSF approach suffered from biased split point selection if few highly predictive variables with few categories were combined with many non-predictive variables with many categories (Scenario 1). However, if covariates with non-linear effects were included (Scenario 2), the linear rank statistic in the CIF approach failed to select these covariates, reducing prediction performance. Of the approximations for MSR-RF, `Lau92` and `HL` performed worse, while the other approximations were comparable.

# 6 Real data example

To demonstrate the efficiency of the proposed approach and to compare the prediction performance with other methods, we analyzed three real datasets on breast cancer with MSR-RF, CIF and RSF. Furthermore, we measured the runtime required for analyzing these datasets.

The `GBSG2` dataset [38] consists of 8 clinical covariates for 686 women from the German breast cancer study group 2, a controlled clinical trial on the treatment of lymph node positive breast cancer for the endpoint recurrence-free survival. The covariates are hormonal therapy (yes or no), age, menopausal status, tumor size and grade, number of positive lymph nodes, progesterone receptor and estrogen receptor. The dataset is freely available in the `pec` R package [32].

The `nki` dataset [39, 40] includes gene expression measurements of 337 lymph node positive breast cancer patients. Seven clinical risk factors and $24,481$ genes are included. We analyzed the endpoint metastasis-free survival. The dataset is freely available in the Bioconductor [41] package `breastCancerNKI` [42].

The `success` study is a randomized phase III study on the treatment of breast cancer, including a genome-wide association study (GWAS). Originally, the dataset included 3322 patients, 17 clinical variables and $693,543$ single nucleotide polymorphisms (SNPs). We performed a standard quality control and linkage disequilibrium pruning ($r^2 > 0.7$), which reduced the dataset to 2781 individuals and $331,195$ SNPs. Finally, to reduce the computational burden and eliminate missing data, we excluded all SNPs with a call fraction below 100%, keeping $151,346$ SNPs. We analyzed the endpoint relapse-free survival. This dataset is available at dbGaP and has ID `phs000547.v1.p1`.

The prediction accuracy was estimated as in Section 5 using 10-fold cross validation and the IBS. The random forest analyses were run using 500 trees and `mtry` $= p/2$, where $p$ is the number of covariates in the model. For MSR-RF, we used $\varepsilon_1 = 0.1$. Because of the size of the datasets, only the minimum

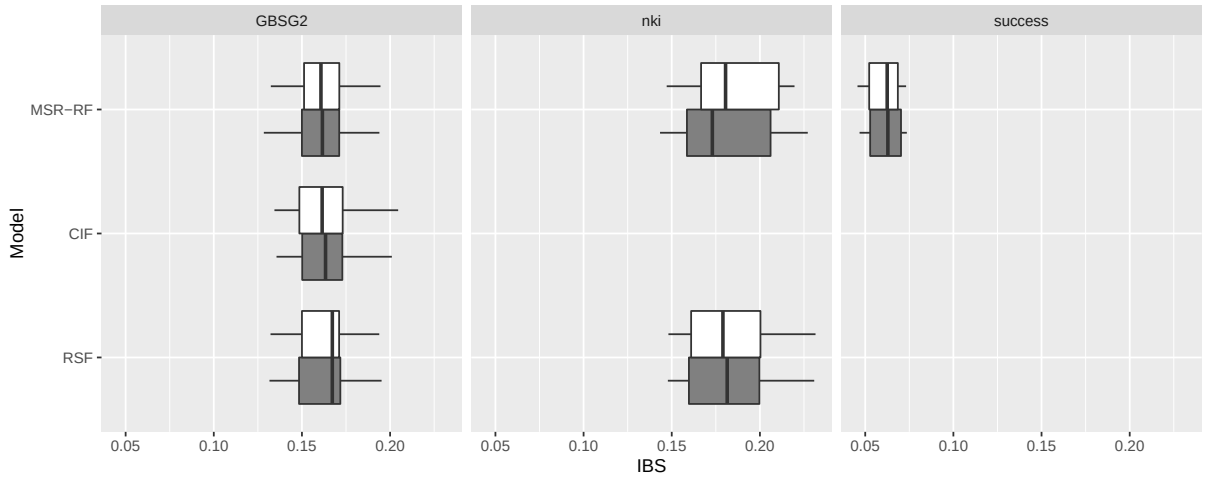


Figure 5: Prediction error for three breast cancer datasets, different analysis methods and resampling schemes. The prediction errors are shown by boxplots (median, quartiles, largest non outliers) of integrated Brier scores (IBS). White boxed correspond to subsampling, where 63.2% of the samples were resampled without replacement. Grey boxes correspond to bootstrapping, where sampling was done with replacement. The maximally selected rank statistics random forest (MSR-RF) results were obtained using the R package `ranger` [37] for the minimum of the $p$-value approximations by Lausen and Schumacher [22] and Lausen *et al.* [23]. Random survival forests [8, RSF] and conditional inference forests [9, CIF] are additionally included.

of the $p$-value approximations by Lausen and Schumacher [22] and Lausen *et al.* [23] was used. All analyses were repeated for sampling with replacement (bootstrapping) and sampling without replacement (subsampling). In each run, all models were trained on the same data. The R packages `randomForestSRC` [14], `party` [20] and `ranger` [37] were used for RSF, CIF and MSR-RF, respectively. The parameters to control the tree size were tuned as in Section 5. The number of trees in the parameter tuning was reduced to 50 caused by runtime restrictions. Tuning results are displayed in Table S2 (supplementary material). Runtime was measured in an additional run without cross validation.

The prediction performance is shown in Figure 5, the runtime in Table 1. In the `GBSG2` dataset, all methods performed similarly. MSR-RF was fastest in computation, followed by RSF and CIF. However, all runtimes for `GBSG2` were very low. In the `nki` dataset, MSR-RF and RSF performed similarly. However, CIF did not complete the analysis. Specifically, we stopped the computation after 4 days without any results. The computation time of MSR-RF was substantially lower than the computation time of RSF. Finally, for the `success` GWAS dataset, only MSR-RF completed the analysis. Both CIF and RSF aborted with a memory error, even if run on a high performance computing node with 192 GB system memory. In all three datasets, no substantial differences were observed between subsampling and bootstrapping. In comparison to the results in Section 5, no outliers were identified, presumably because the number of trees grown per random forest was increased.

Table 1: Runtime (in minutes) to grow a random forest with 500 trees using the new maximally selected rank statistics random forest (MSR-RF) approach, conditional inference forests [9, CIF] and random survival forests [8, RSF]. The MSR-RF results were obtained using the R package `ranger` [37] and the minimum of the $p$-value approximations by Lausen and Schumacher [22] and Lausen *et al.* [23]. All results are for sampling with replacement and the tuned parameters. For MSR-RF and RSF, multithreading was used with 12 threads. *Stopped after 4 days of computation. †Memory error.

| Method | Runtime in minutes | | |
|---|---|---|---|
| | `GBSG2` | `nki` | `success` |
| MSR-RF | 0.01 | 2.23 | 48.17 |
| CIF | 0.05 | $> 5760^*$ | NA† |
| RSF | 0.04 | 26.37 | NA† |

# 7 Discussion

Random forests for survival data with maximally selected rank statistics (MSR-RF) reduce split point selection bias and are able to deal with non-linearity in the covariates when compared with RSF and CIF. RSF is based on the standard random forest algorithm, and it showed the expected split point selection bias [17] in Monte-Carlo simulation studies. Specifically, MSR-RF and CIF both outperformed RSF in the simulation scenarios with mixed covariate types because RSF favors covariates with many possible split points over covariates with fewer possible split points. In a simulation study with non-linear covariate effects, we observed a better performance of MSR-RF and RSF when compared with CIF. This finding can be explained by the fact that the standard CIF utilizes a linear rank statistic, which cannot cope with non-linearities. However, we stress that none of the random forest methods failed completely on any of the considered simulated or real data sets. In fact, the random forest methods performed equally well of even better than the standard Cox model in our simulations.

With real data examples, we have demonstrated the computational efficiency of the proposed MSR-RF approach. In a dataset with few clinical covariates it performed as well as CIF and RSF, and all runtimes were low. In a gene expression dataset, CIF was unable to complete the analysis, while MSR-RF and RSF had comparable prediction errors. However, MSR-RF was substantially faster than RSF in this example. Finally, in a genome-wide association study, the analysis could be completed with the fast implementation of MSR-RF only.

However, the MSR-RF version implemented in `ranger` is a compromise between the reduction in split point selection bias and computational efficiency. In fact, in our comparison of 6 approaches for approximating the maximally selected rank statistic, only the Monte-Carlo permutation approach and the maximally selected Gauss statistic achieved unbiased split point selections. However, both approaches are computationally intensive which restricts their usage to small random forests. Of the remaining methods, the minimum of the approximations by Lausen *et al.* [23] and Lausen and Schumacher [22] achieved the best results. The method performed equally well for small sample sizes. Since the minimum approach is fast to compute, we consider this approach as the best compromise in the tradeoff between split point selection bias and runtime for larger random forests.

The proposed method can be embedded in the conditional inference framework for recursive partitioning by Hothorn *et al.* [9], which is the basis for CIF. Using the notation from Hothorn *et al.* [9], the influence function $h(\cdot)$ can be chosen to represent the log-rank scores $a_i$, as described in Section 2. The covariate transformation $g_j(\cdot)$ would need to be specified as $g_j(x_{ji}) = \sum_{k=r(x_{ji})}^{n} e_n(k)$, where $r(x_{ji})$ is the rank of the observation $i$ in the covariate $j$ under the assumption of sorted covariates and no ties, and $e_n(k)$ is the unit vector of length $n$ with the $k$th element equal to one [9]. In principle, this could be implemented in the `party` package [20]. However, the $p$-value approximations described in Section 2 would not be available. Furthermore, the determination of the best split point in the selected covariate would not be required in the second step, since it is already known from the first step.

MSR-RF are just one special case of the general approach to maximally selected statistics [43]. It

is easily extenable to continuous outcome by setting the scores equal to the ranks [22]. Furthermore, the approach can be extended to categorical endpoints by using a maximally selected $\chi^2$ statistic [44]. For the binary case, exact methods to obtain $p$-values are known [45–47], and for the categorical case approximations have already been proposed [48]. These approaches could be utilized to implement classification and regression random forests using maximally selected statistics.

The restriction of possible split points to an interval $[\mu_1, \mu_2]$, corresponding to quantiles $\varepsilon_1, \varepsilon_2$ of the distribution of the covariates, suppresses extreme end-cut splits, to provide sufficient sample sizes in both groups for asymptotic properties to hold for the $p$-value approximation [22]. However, recent findings [49] indicate that the ability to produce end-cut splits is desirable in random forests and our results in Section 4 suggest that the approximations work well for small sample sizes. Therefore, one could use a smaller value for the parameter `minprop` or even set it to 0, not restricting the possible split points at all and possibly increasing the prediction accuracy in high-noise scenarios.

In this work, we used the Benjamini and Hochberg [30] approach to adjust for multiple testing and to decide whether tree growing should be stopped. Other approaches could, of course, be used to adjust for the multiple testing of `mtry` variables at a node. A natural choice would be Hommel's procedure [50] which relies on the positive stochastic dependence of covariates [51]. However, we consider these alternative approaches to be computationally more demanding. Computational efficiency has therefore led to the choice of the Benjamini Hochberg procedure. The package `ranger` is quite flexible, and the interested user can easily add his/her preferred procedure.

Our results suggest that the choice of the resampling scheme, i.e. subsampling or bootstrapping, has no general impact on the prediction performance. However, the tuned parameters were different, indicating that for the same parameter combination the prediction performance of the methods differs. Since unbiased split point selection is only possible with subsampling (Section 4) and most of the theoretical results on random forests are obtained in this case [13, 52], we generally recommend to use subsampling.

Random forests with maximally selected rank statistics (MSR-RF) are a useful alternative to other random forest approaches for survival analysis. By the use of a two-step procedure, bias in split point selection is reduced considerably. The method outperforms the Cox model on simulated datasets. Compared to other random forest approaches for survival analysis, the method performs better than random survival forests if informative dichotomous variables are combined with uninformative variables with more categories and better than conditional inference forests if non-linear covariate effects are included. In none of the analyzed scenarios the method is outperformed by another approach. With simple $p$-value approximations the method is computationally efficient and capable of analyzing very large datasets.

## Acknowledgements

This work was financially supported by the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement No. HEALTH-F2-2011-278913 (BiomarCaRE).

## References

[1] Cox DR. Regression models and life-tables. *Journal of the Royal Statistical Society. Series B (Methodological)* 1972; **34**(2):187–220.

[2] Tibshirani R. The lasso method for variable selection in the Cox model. *Statistics in Medicine* 1997; **16**(4):385–395.

[3] Breiman L. Random forests. *Machine Learning* 2001; **45**(1):5–32.

[4] Davis RB, Anderson JR. Exponential survival trees. *Statistics in Medicine* 1989; **8**(8):947–961.

[5] Hothorn T, Lausen B, Benner A, Radespiel-Tröger M. Bagging survival trees. *Statistics in Medicine* 2004; **23**(1):77–91.

[6] Keleş S, Segal MR. Residual-based tree-structured survival analysis. *Statistics in Medicine* 2002; **21**(2):313–326.

[7] Bou-Hamad I, Larocque D, Ben-Ameur H. A review of survival trees. *Statistics Surveys* 2011; **5**:44–71.

[8] Ishwaran H, Kogalur UB, Blackstone EH, Lauer MS. Random survival forests. *The Annals of Applied Statistics* 2008; **2**(3):841–860.

[9] Hothorn T, Hornik K, Zeileis A. Unbiased recursive partitioning: A conditional inference framework. *Journal of Computational and Graphical Statistics* 2006; **15**(3):651–674.

[10] Scornet E, Biau G, Vert JP. Consistency of random forests. *The Annals of Statistics* 2015; **43**(4):1716–1741.

[11] Ishwaran H, Kogalur UB. Consistency of random survival forests. *Statistics & Probability Letters* 2010; **80**(13–14):1056–1064.

[12] Wager S, Walther G. Uniform convergence of random forests via adaptive concentration. *Preprint arXiv:1503.06388*, 2015.

[13] Wager S. Asymptotic theory for random forests. *Preprint arXiv:1405.0352*, 2014.

[14] Ishwaran H, Kogalur UB. randomForestSRC: Random forests for survival, regression and classification. R package version 1.5.5, 2014.

[15] Boulesteix AL, Janitza S, Hapfelmeier A, Van Steen K, Strobl C. Letter to the Editor: On the term 'interaction' and related phrases in the literature on random forests. *Briefings in Bioinformatics* 2015; **16**(2):338–345.

[16] Loh WY. Fifty years of classification and regression trees. *International Statistical Review* 2014; **82**(3):329–348.

[17] Strobl C, Boulesteix AL, Zeileis A, Hothorn T. Bias in random forest variable importance measures: Illustrations, sources and a solution. *BMC Bioinformatics* 2007; **8**(1):25.

[18] Ziegler A, König IR. Mining data with random forests: current options for real-world applications. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 2014; **4**(1):55–63.

[19] Yang P, Hwa Yang Y, B Zhou B, Y Zomaya A. A review of ensemble methods in bioinformatics. *Current Bioinformatics* 2010; **5**(4):296–308.

[20] Hothorn T, Hornik K, Strobl C, Zeileis A. party: A laboratory for recursive partytioning. R package version 1.0-19, 2014.

[21] Strobl C. Discussion: Fifty years of classification and regression trees. *International Statistical Review* 2014; **82**(3):349–352.

[22] Lausen B, Schumacher M. Maximally selected rank statistics. *Biometrics* 1992; **48**(1):73–85.

[23] Lausen B, Sauerbrei W, Schumacher M. Classification and regression trees (CART) used for the exploration of prognostic factors measured on different scales. *Computational Statistics*, Dirschedl P, Ostermann R (eds.). Physica Verlag, Heidelberg, 1994; 483–496.

[24] Lausen B, Hothorn T, Bretz F, Schumacher M. Assessment of optimal selected prognostic factors. *Biometrical Journal* 2004; **46**(3):364–374.

[25] Hothorn T, Lausen B. On the exact distribution of maximally selected rank statistics. *Computational Statistics & Data Analysis* 2003; **43**(2):121–137.

[26] Worsley KJ. An improved Bonferroni inequality and applications. *Biometrika* 1982; **69**(2):297–302.

[27] Streitberg B, Röhmel J. Exact distributions for permutation and rank tests: An introduction to some recently published algorithms. *Statistical Software Newsletter* 1986; **12**(1):10–17.

[28] Hájek J, Šidák Z, Sen PK. *Theory of Rank Tests*. 2 edn., Academic Press, New York, 1999.

[29] Genz A. Numerical computation of multivariate normal probabilities. *Journal of Computational and Graphical Statistics* 1992; **1**(2):141–149.

[30] Benjamini Y, Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)* 1995; **57**(1):289–300.

[31] Graf E, Schmoor C, Sauerbrei W, Schumacher M. Assessment and comparison of prognostic classification schemes for survival data. *Statistics in Medicine* 1999; **18**(17-18):2529–2545.

[32] Mogensen UB, Ishwaran H, Gerds TA. Evaluating random forests for survival analysis using prediction error curves. *Journal of Statistical Software* 2012; **50**(11):1–23.

[33] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing: Vienna, Austria, 2014.

[34] Wickham H. *Advanced R*. CRC Press, Boca Raton, 2014.

[35] Peters A, Hothorn T. ipred: Improved predictors. R package version 0.9-3, 2013.

[36] Hothorn T. maxstat: Maximally selected rank statistics. R package version 0.7-21, 2015.

[37] Wright MN, Ziegler A. ranger: A fast implementation of random forests for high dimensional data in C++ and R. *Journal of Statistical Software* 2016; In press.

[38] Schumacher M, Bastert G, Bojar H, Hübner K, Olschewski M, Sauerbrei W, Schmoor C, Beyerle C, Neumann RL, Rauschecker HF. Randomized 2 x 2 trial evaluating hormonal treatment and the duration of chemotherapy in node-positive breast cancer patients. German Breast Cancer Study Group. *Journal of Clinical Oncology* 1994; **12**(10):2086–2093.

[39] Van De Vijver MJ, He YD, van't Veer LJ, Dai H, Hart AA, Voskuil DW, Schreiber GJ, Peterse JL, Roberts C, Marton MJ, *et al.*. A gene-expression signature as a predictor of survival in breast cancer. *New England Journal of Medicine* 2002; **347**(25):1999–2009.

[40] van 't Veer LJ, Dai H, van de Vijver MJ, He YD, Hart AAM, Mao M, Peterse HL, van der Kooy K, Marton MJ, Witteveen AT, *et al.*. Gene expression profiling predicts clinical outcome of breast cancer. *Nature* 2002; **415**(6871):530–536.

[41] Gentleman RC, Carey VJ, Bates DM, Bolstad B, Dettling M, Dudoit S, Ellis B, Gautier L, Ge Y, Gentry J, *et al.*. Bioconductor: open software development for computational biology and bioinformatics. *Genome Biology* 2004; **5**(10):R80.

[42] Schroeder M, Haibe-Kains B, Culhane A, Sotiriou C, Bontempi G, Quackenbush J. breastCancerNKI: Genexpression dataset published by van't Veer et al. 2002 and van de Vijver et al. 2002 (NKI). R package version 1.3.1, 2011.

[43] Hothorn T, Zeileis A. Generalized maximally selected statistics. *Biometrics* 2008; **64**(4):1263–1269.

[44] Miller R, Siegmund D. Maximally selected chi square statistics. *Biometrics* 1982; **38**:1011–1016.

[45] Koziol JA. On maximally selected chi-square statistics. *Biometrics* 1991; **47**:1557–1561.

[46] Boulesteix AL. Maximally selected chi-square statistics and binary splits of nominal variables. *Biometrical Journal* 2006; **48**(5):838–848.

[47] Boulesteix AL. Maximally selected chi-square statistics for ordinal variables. *Biometrical Journal* 2006; **48**(3):451–462.

[48] Betensky RA, Rabinowitz D. Maximally selected x2 statistics for kx2 tables. *Biometrics* 1999; **55**(1):317–320.

[49] Ishwaran H. The effect of splitting on random forests. *Machine Learning* 2015; **99**(1):75–118.

[50] Hommel G. A stagewise rejective multiple test procedure based on a modified Bonferroni test. *Biometrika* 1988; **75**(2):383–386.

[51] Goeman JJ, Solari A. Multiple hypothesis testing in genomics. *Statistics in Medicine* 2014; **33**(11):1946–1978.

[52] Biau G, Scornet E. A random forest guided tour. *Preprint arXiv:1511.05741*, 2015.